
Microsoft
WinHEC
2005
Microsoft
Windows
Hardware & Driver Central

x86 Everywhere

Chris Herring
Director Strategy
PCSG
AMD

Session Outline

- Instructions, Languages, and the Euro
- Architectural Evolution
  - Cost of Deployment
  - Macro Level
- Extending x86
  - Architectures du Jour
  - Evolution/Convergence: Why Now?
- Primary Drivers: Performance, Power, Cost
  - Knobs and Levers
  - Absolute Costs and Tradeoffs
- Future Innovations

Session Goals

- An understanding of the benefits of a unified approach to Instruction Set Architecture
- No limitations to extending x86 architecture
- The time is now
- This is not a crazy idea!

Imagine if instruction sets were languages…

Bonjour Monde
Hallo Welt
Γειάσου κόσμος
Ciao Mondo
□□□□□□□
□□□□ □□
Hello Mundo
Здравствулте! Мир
Hola Mundo

Or, conversely, if languages were instruction set architectures…

Language Should Unify

- What I really want to say is “Hello, World”
- Each Instruction Set Architecture (ISA) requires translation
- Translation, Virtual Machines, Transmorphing, Transcoding, Simulation, Re-compile: All “tech speak” for 2nd language communication
- What we have in effect is, truly, an ISA “Tower of Babel”

SHARC®

UltraSPARC®

XScale™

PowerPC™

SuperH®

SIGMA DESIGNS®

Cell

Towards Instruction Set Consolidation

Future innovation should come in micro-architecture enhancements and compatible extensions to dominant instruction sets, rather than the creation of new instruction sets.

1

Is the trend clear?

With ever growing software complexity and installed base the value of remaining compatible with and extending existing, dominant instruction sets heavily outweighs any disadvantages.

2

Is the time now?

Technology has passed the point where instruction set costs are no longer relevant.

What The Euro Can Teach Us

The economic benefits of moving away from multiple currencies is enormous.

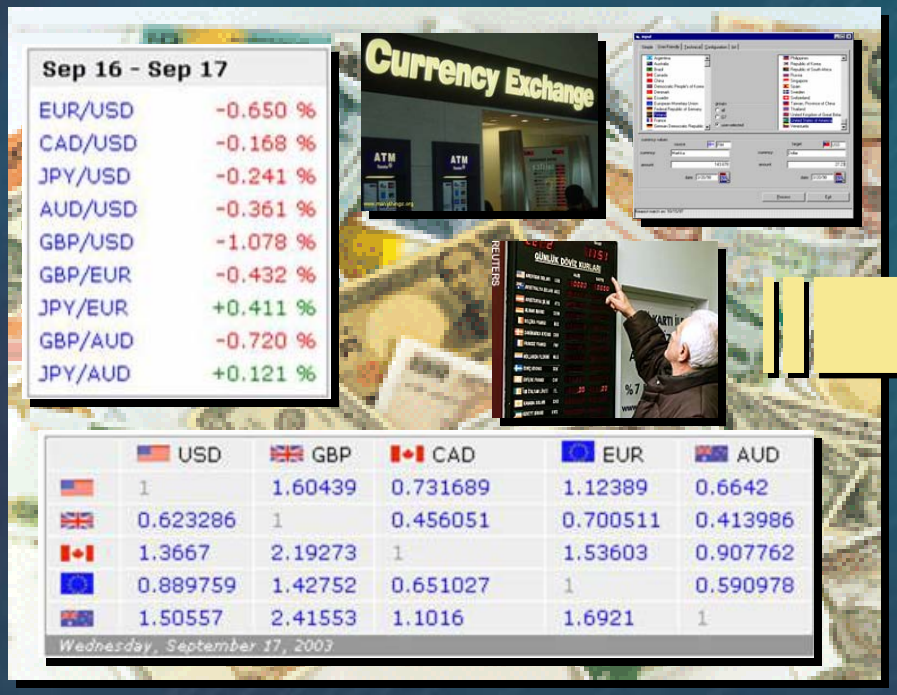

$36 billion **per year in avoided transaction costs, or** ***$90 per EU resident***

What The Euro Can Teach Us

What were the catalysts that prompted standardization?

- ✓ **Growing markets**
- ✓ **More sophisticated consumers**
- ✓ **Desire for increased stability**

Architectural Evolution
Macro-Level

70s 80s 90s 00s 10s

Growing Software Complexity

Majority of System Software and Applications are x86-Based

Time is Right for Full-Function x86 on All Form Factors

Exploding Functionality per Square Inch

The True Cost of Deployment

Possible Solutions

Port thousands of applications, operating systems, drivers, codecs, tool chains and virtual machines.

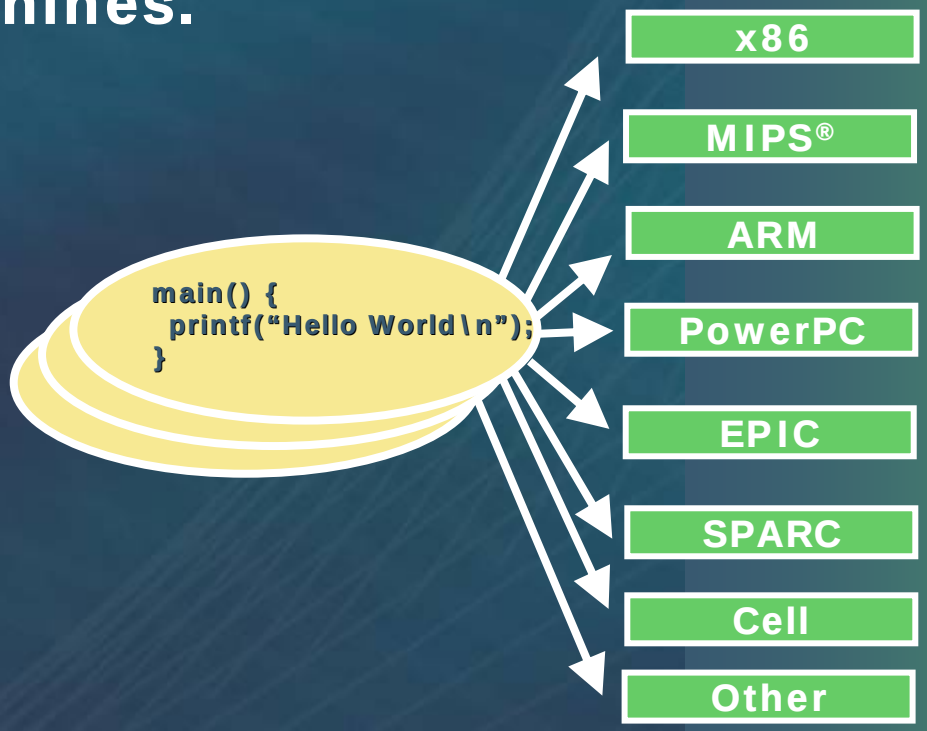

- ***Resource intensive***
- ***Slow***
- ***Costly to maintain***

Possible Solutions

Develop a web-based interface for each application.

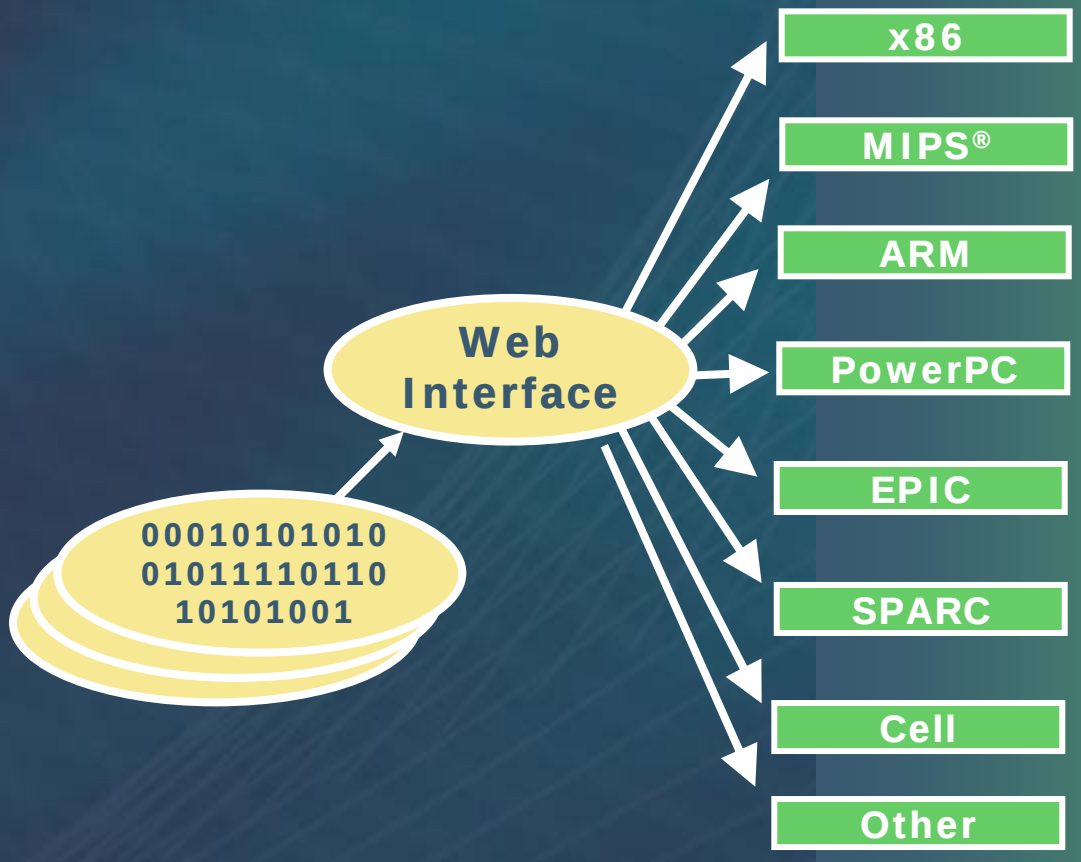

- ***Assumes always-connected client server***
- ***Limited functionality***
- ***Least Common Denominator***
- ***Difficult security***

Possible Solutions

Write once (or port once) and run anywhere.

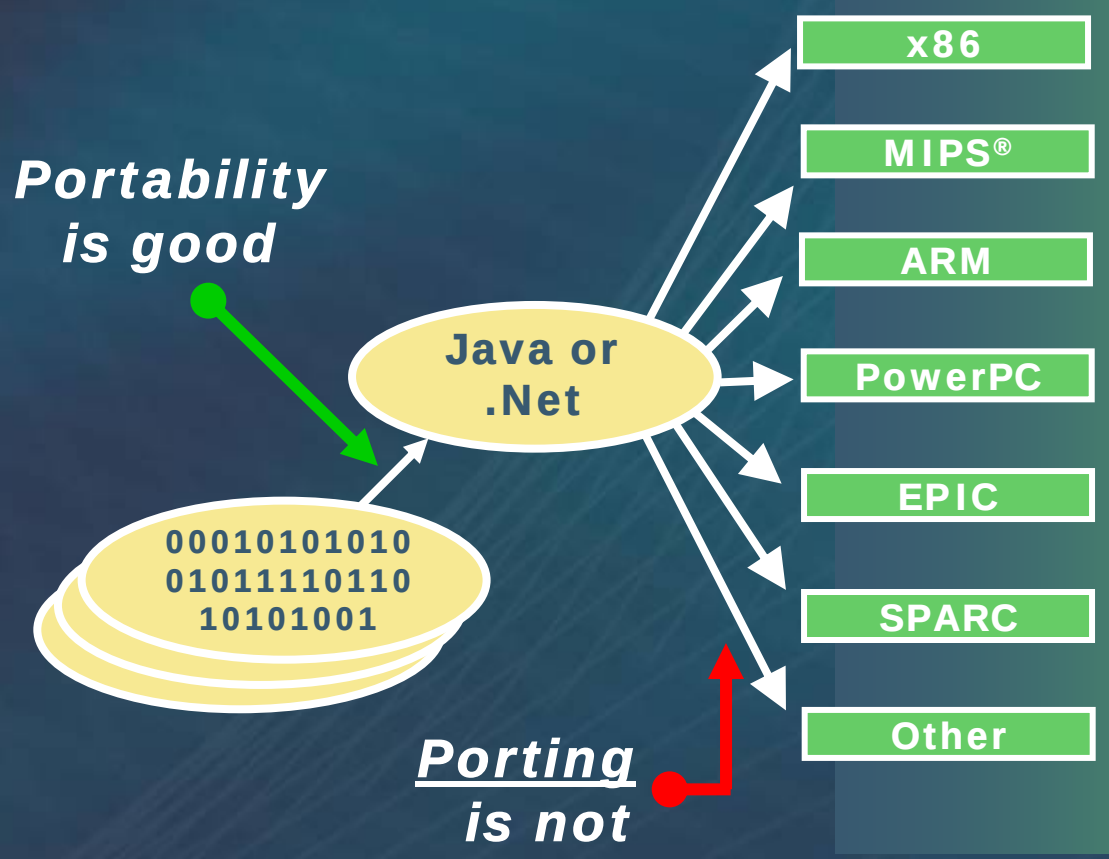

YIELD

- *Heavy testing*
- *Ongoing optimization*
- *Per-platform customization*

Is it really “port once”

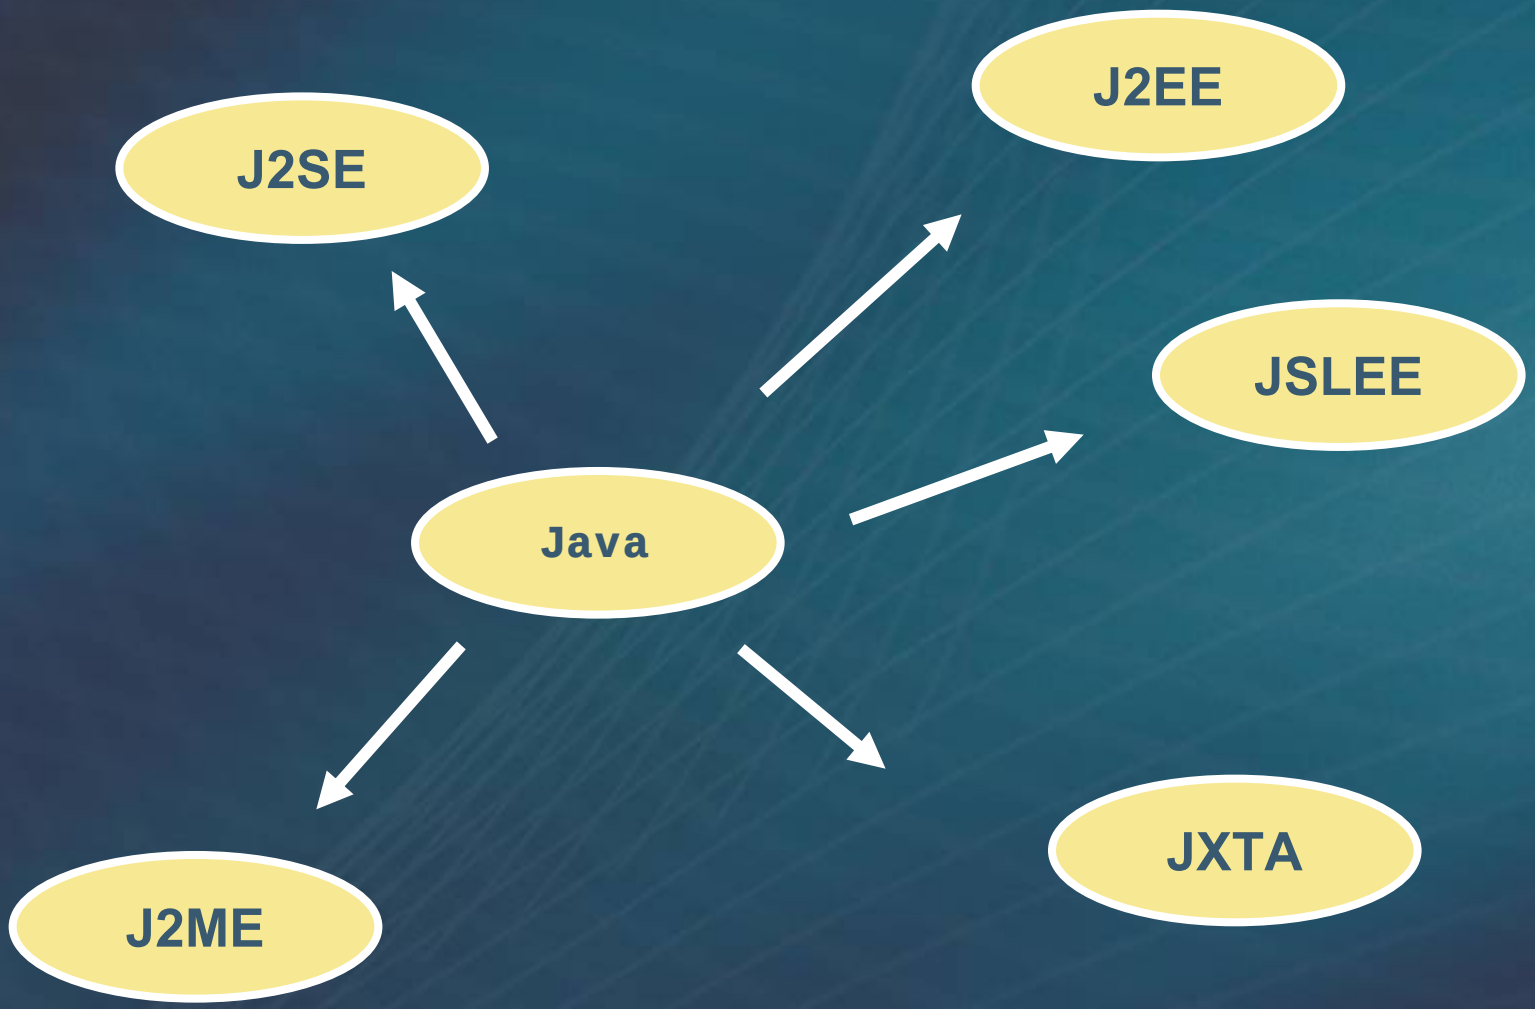

Executing and/or translating to multiple languages and platforms is a necessary cost — not something to be desired.

Architectural Evolution Macro-Level

Common Instruction Set Architecture

- No need to port
- No need for multiple validations
- Built in OS integration
- Robust security
- Investment protection

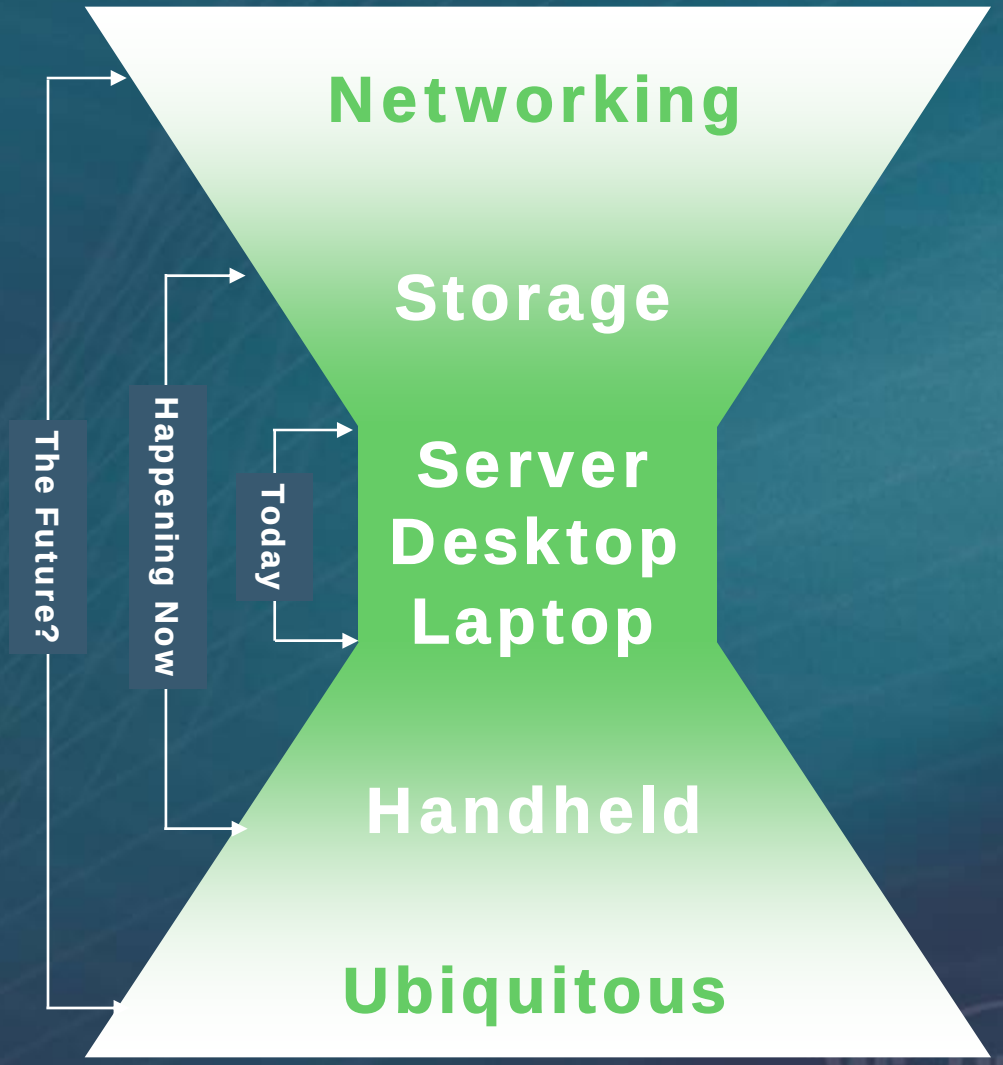

The Importance of Dusty Decks… Or Data Sets Never Die…

Essential legacy data and code lives forever.

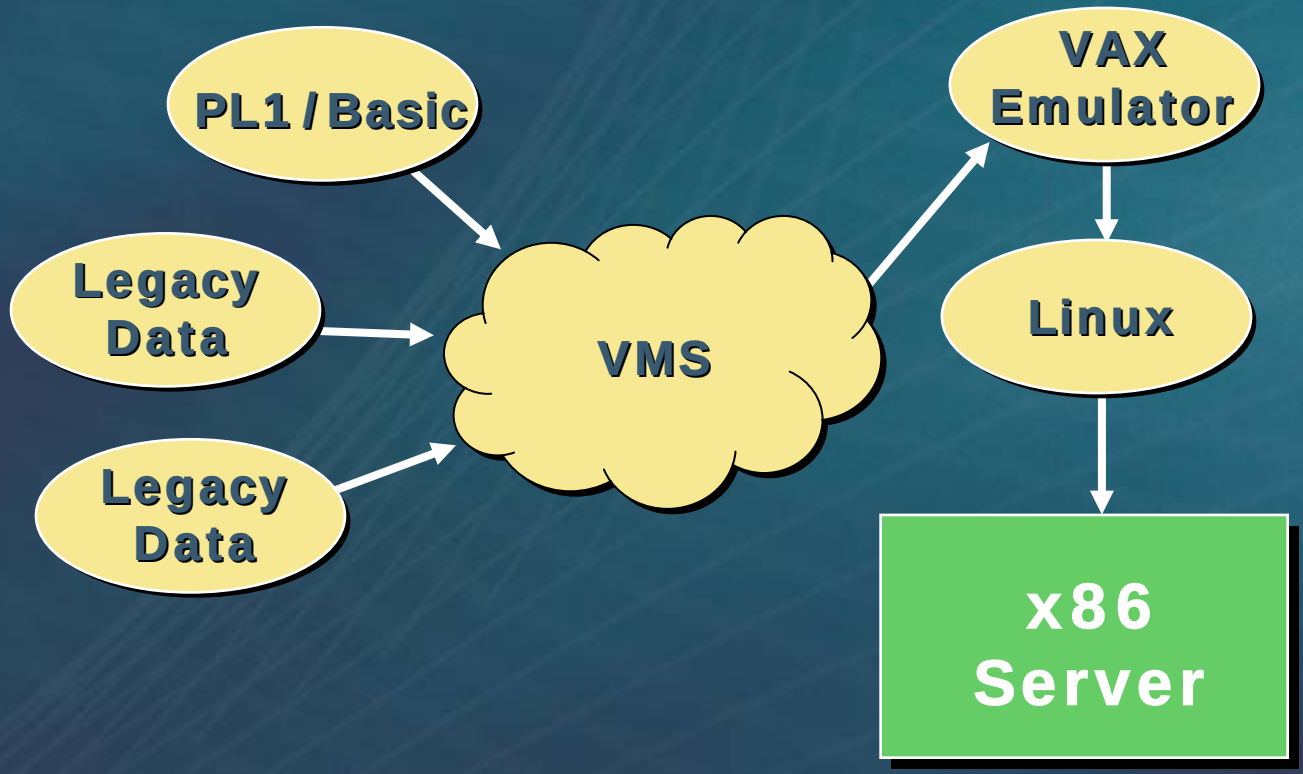

Extending x86

Enterprise x86

Consumer x86

SAN / NAS

Rugged Small Form Factor

Internet Appliances

Low / Mid Server

LFF Consumer Electronics

Handheld

High End Server

Networking

Workstation

Desktop

SFF Consumer Electronics

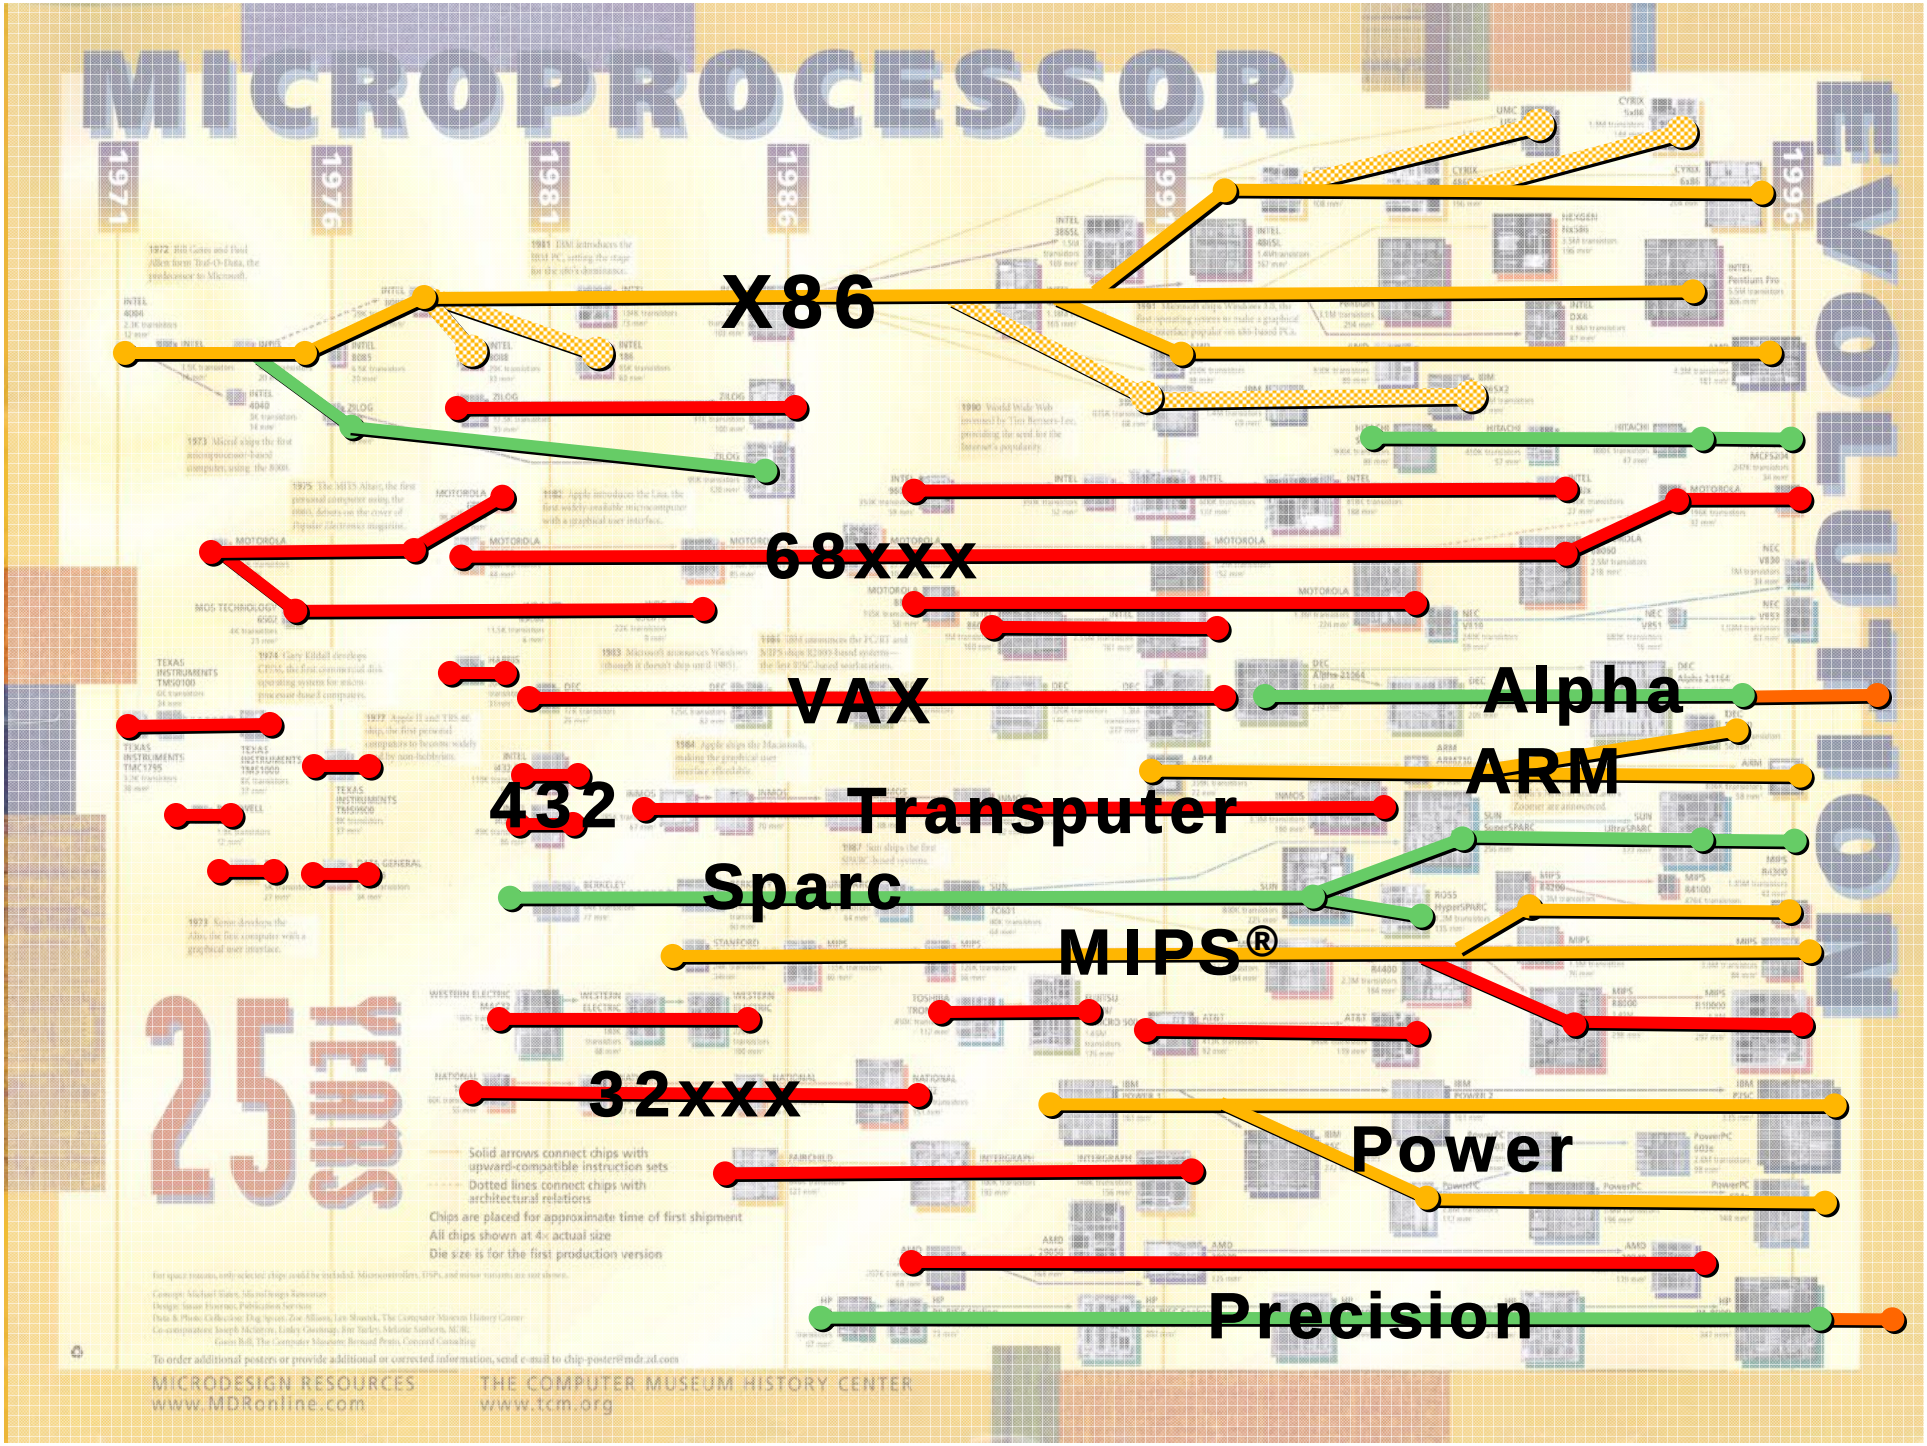
MICROPROCESSOR
EVOLUTION
X86
68xxx
VAX
Alpha
ARM
432
Transputer
Sparc
MIPS®
32xxx
Power
Precision
25 YEARS
Solid arrows connect chips with upward-compatible instruction sets
Dotted lines connect chips with architectural relations
Chips are placed for approximate time of first shipment
All chips shown at 4× actual size
Die size is for the first production version
MICRODESIGN RESOURCES
www.MDRonline.com
THE COMPUTER MUSEUM HISTORY CENTER
www.tcm.org

ISA Architectural Evolution/ Convergence (architectures must cross platform boundaries)

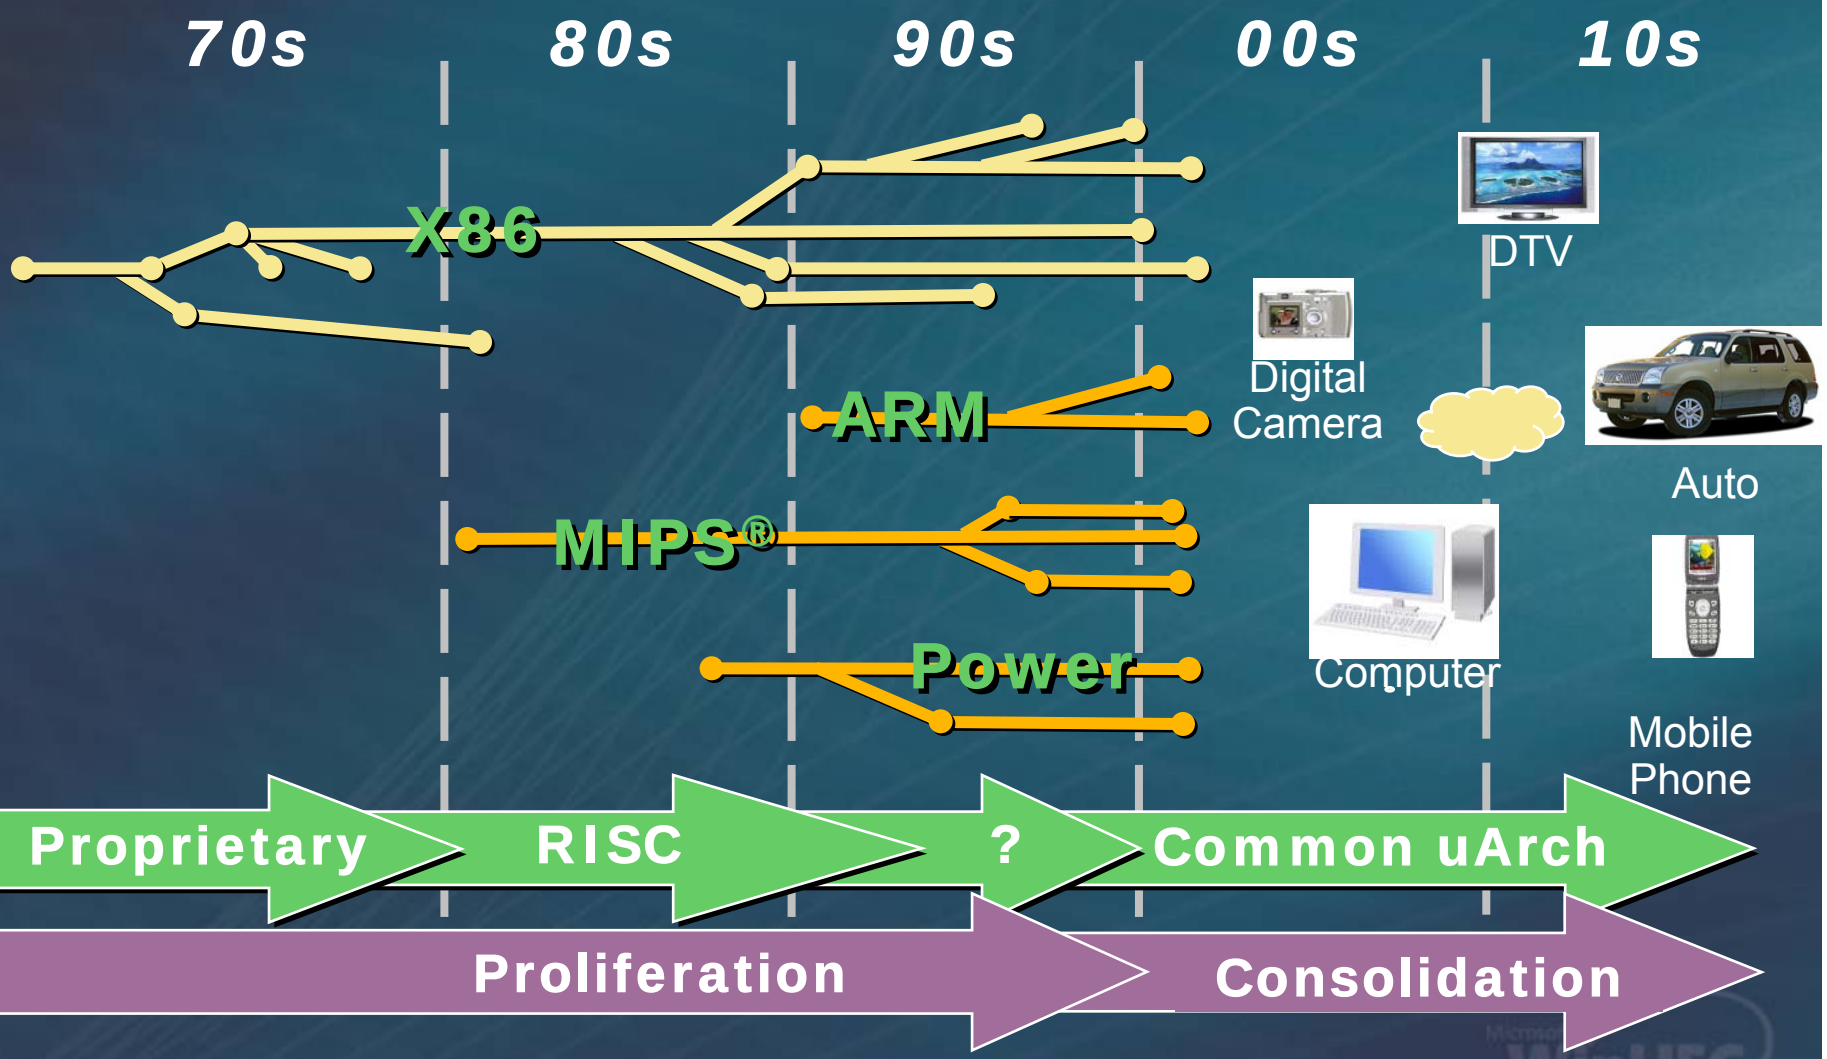

Why is It Possible Now?

- Moore's Law
  - Core processors can now be so small that any overhead of x86 is easily affordable (a few mm², a few % of total die, trivial power increment)
- Sufficiency of performance
  - CPU designs can range from small simple designs to huge server designs
- Added functionality makes processor core small part of chip
  - Large L1 and L2 cache
  - SOC functions (memory control, graphics, etc.)
- Pads fundamentally force minimum die size much larger than core
- Learning
  - Lots of design tricks have accumulated

Interconnect Evolution

AMD Am386® Processor
2 Alum. layer
1.0µ design rules

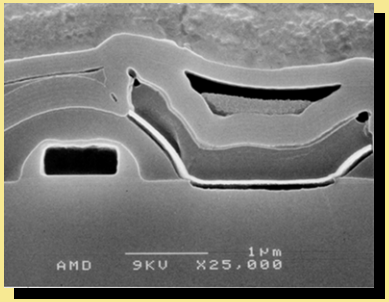

AMD Athlon™ Processor
6 Copper layer
180 nm design rules

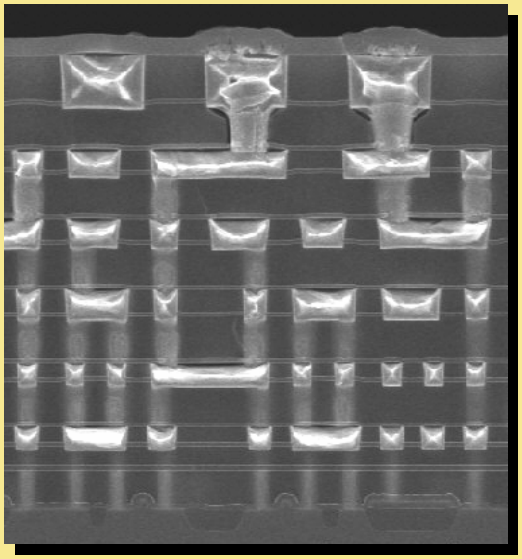

AMD Athlon™ 64 Processor
9 Copper layer
130 nm design rules

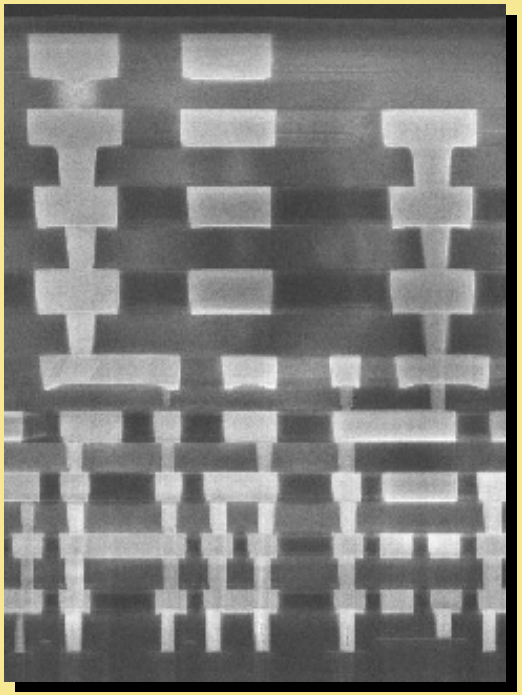

Planned Transistor Evolution

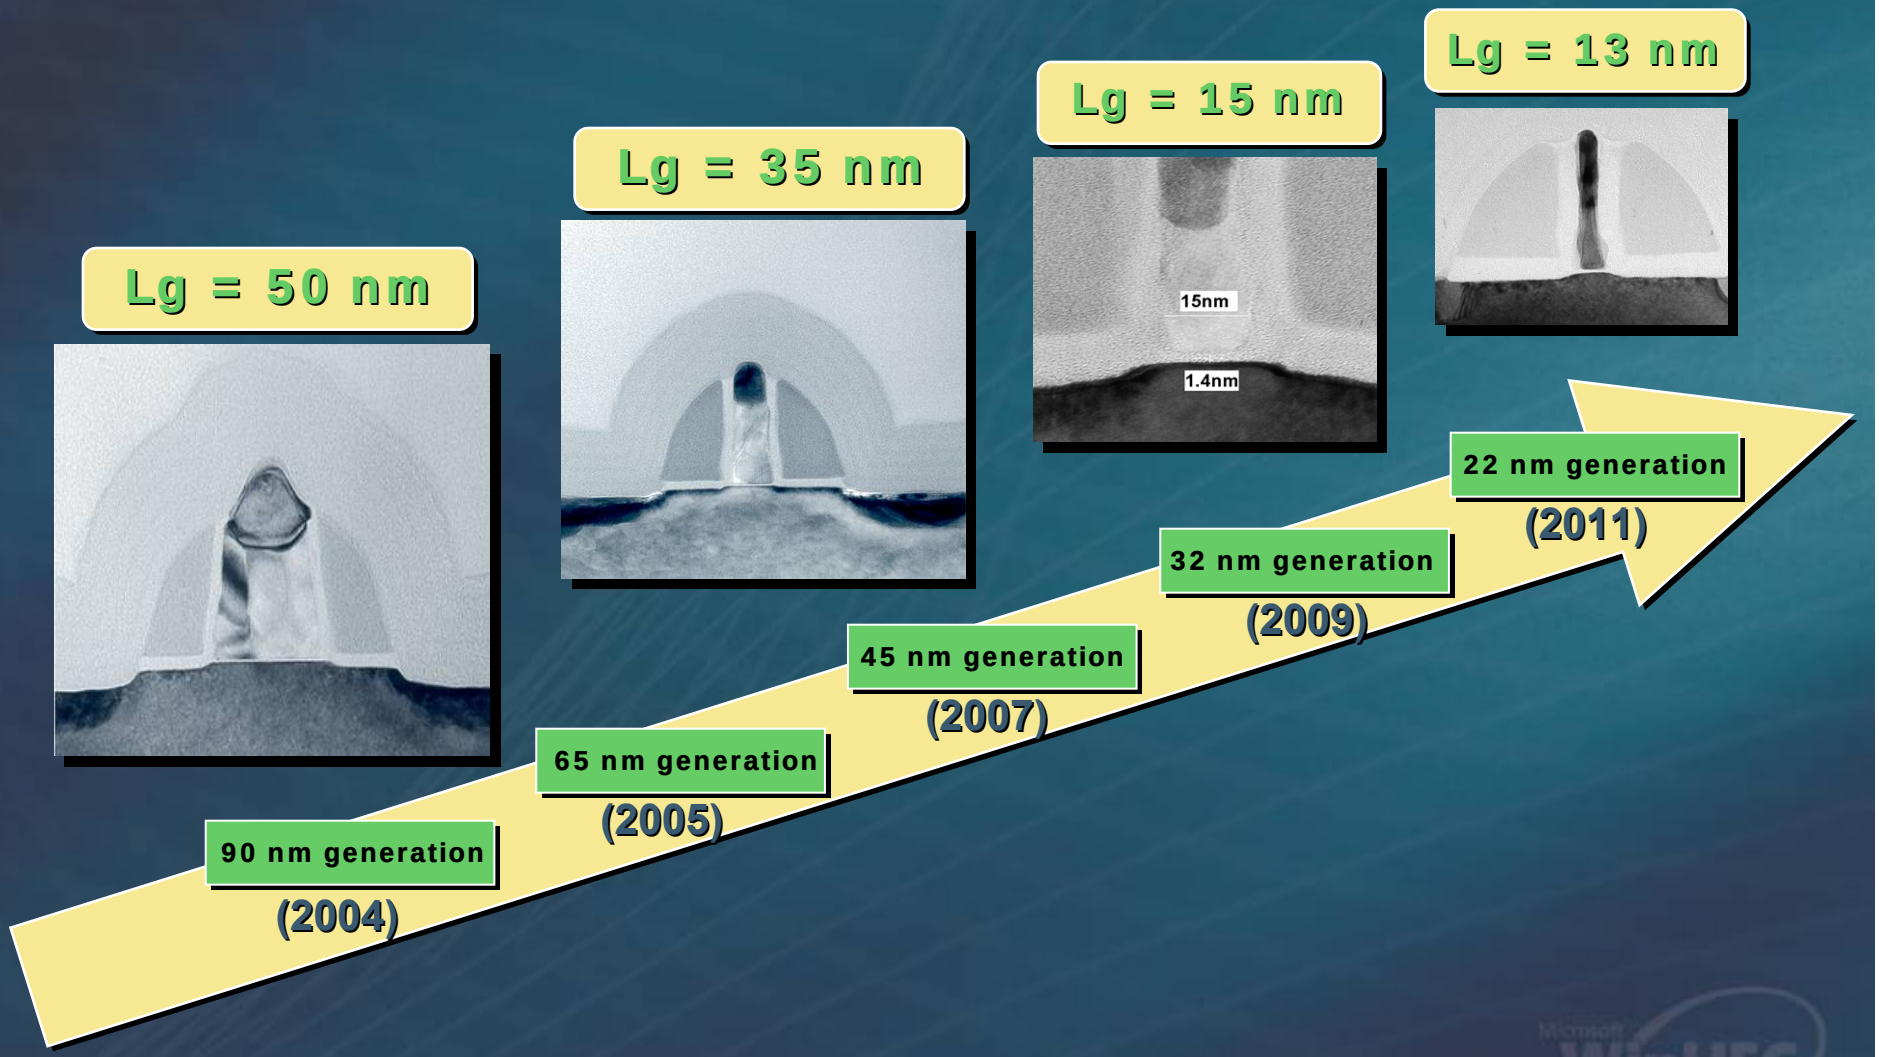

Huge Variation – Without Even Trying!

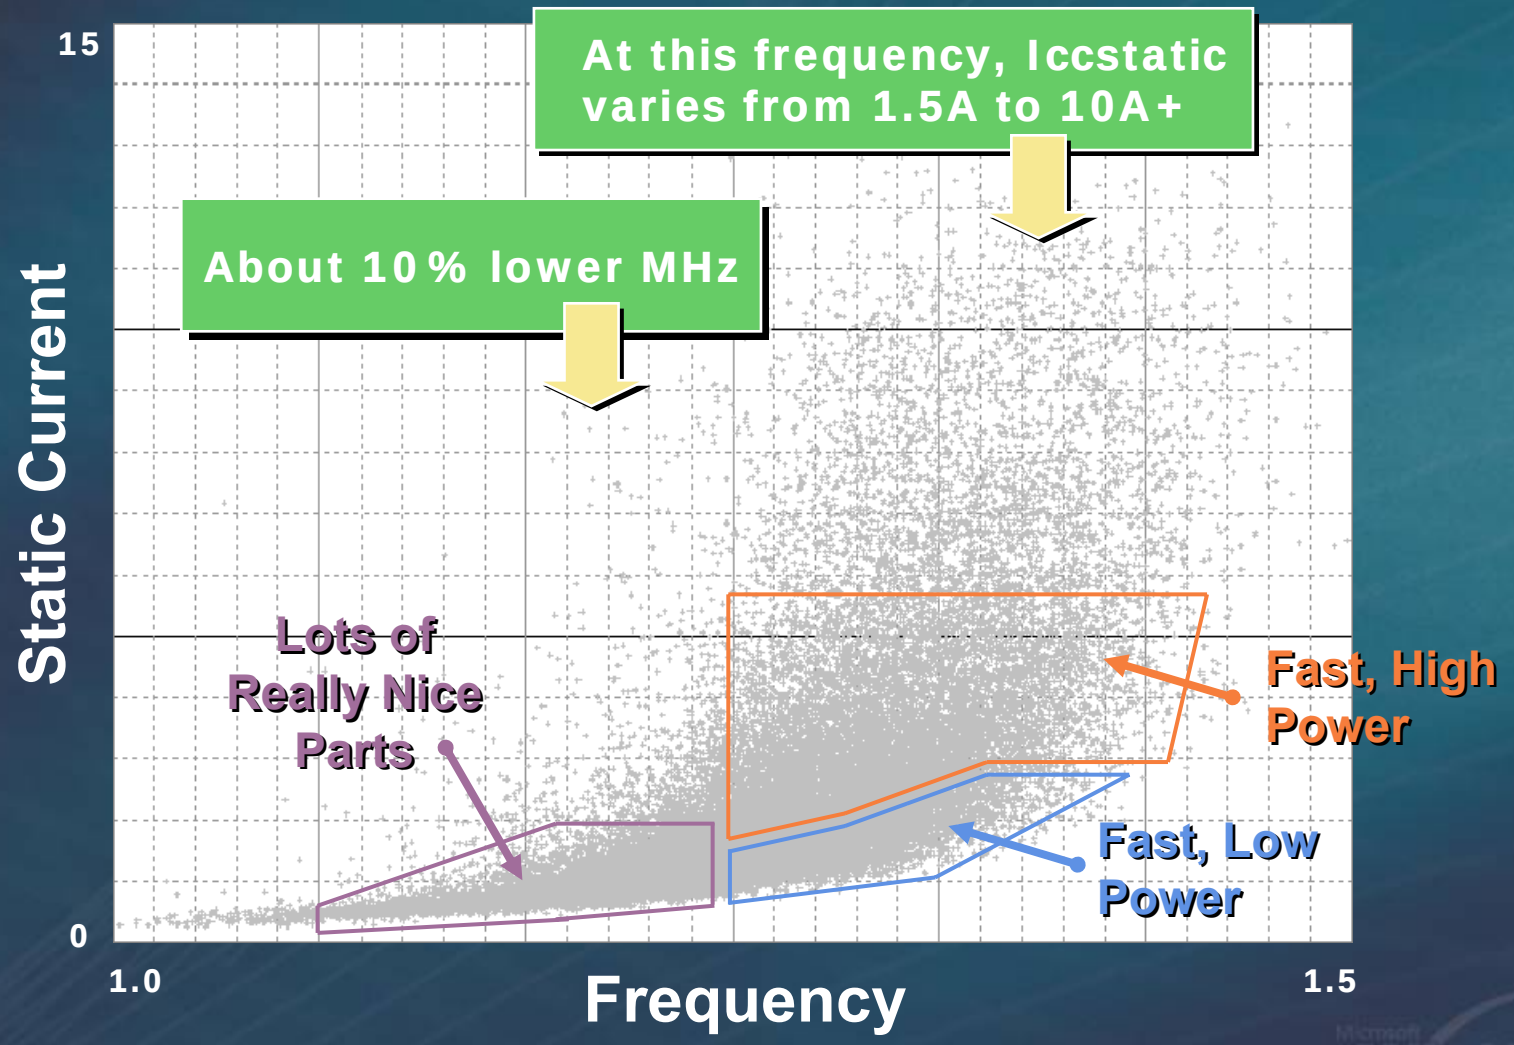

Primary Drivers of Performance, Power and Cost

Performance

Memory Latency — We Have Hit the Memory Wall

- At 1GHz and 1 IPC a 1% effective miss rate cuts performance by approx. 50%, 4% cuts it by approx. 80%
- Required IPC, application footprint and DRAM speed dictate cache size

Peak Instruction Rate

- Determined by IPC and MHz

Power

Dynamic

- Voltage
- Switching capacitance
  - # of parallel units
  - # of flops (pipeline depth)
  - Clock skew goal
  - Long buses
  - Clock gating
  - I/O

Static

- Process: Gate leakage, Dcap leakage and Ioff
- Voltage and temperature
- Total transistor width
  - Leaky transistor width if substantial use of low leakage transistors

Cost

Package, Assembly and Test

- 25-50% of total cost

Cache and IO Often Dominate Die Size

- IO often in the range of 15-20% of die area
- Cache can be as much as 50% of the die area or more

Dynamic Range of Design Choices

Lever	Range	Factor
Pipeline Depth	5-25 stages	5x
Cache	16K-2M bytes	125x
ILP	1-10 units	10x
Voltage	0.7-1.8 volts	2.5x
IO Pins	100-500	5x
Frequency	300-3000 MHz	10x
Leakage (Idoff)	1-1000 nA/um	1000x
Die Size	20-300 mm²	15x
Dynamic Power	0.2-100 W	500x

All can be varied independent of instruction set.

Three Orders of Magnitude Transistor Count

16	→ 32 →	64 Bits
Integer	→	FPU, SIMD, Vector
1/3 Issue	→	9 Issue
Trivial Cache	→	1MB Cache
CPU	→	SOC

AMD K6®-III Processor

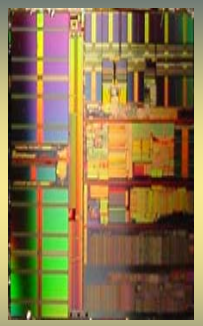

0.25µ
9M transistors
78mm²

AMD Athlon™ Processor

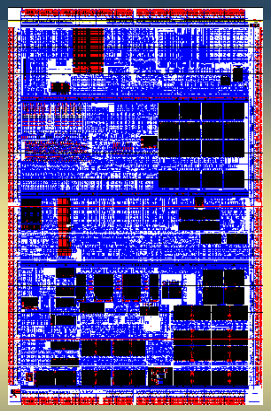

0.18µ
37M transistors
120mm²

AMD Opteron™ Processor

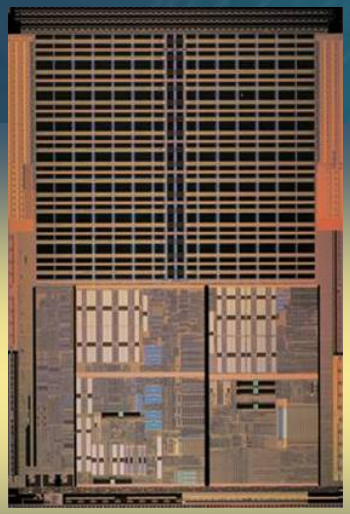

0.13µ
100M transistors
193mm²

AMD AMD64™ Processor

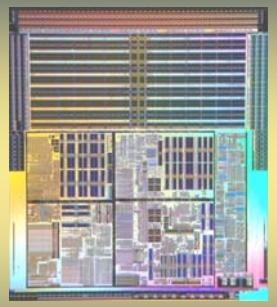

90nm
84mm²

AMD Geode™ Processor

0.15µ
9.5M transistors
58.3mm²

Future AMD Geode™ SOC Processor

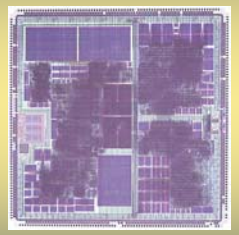

90nm
30M+ transistors
40-45mm²

Levers Affecting Performance, Power and Cost

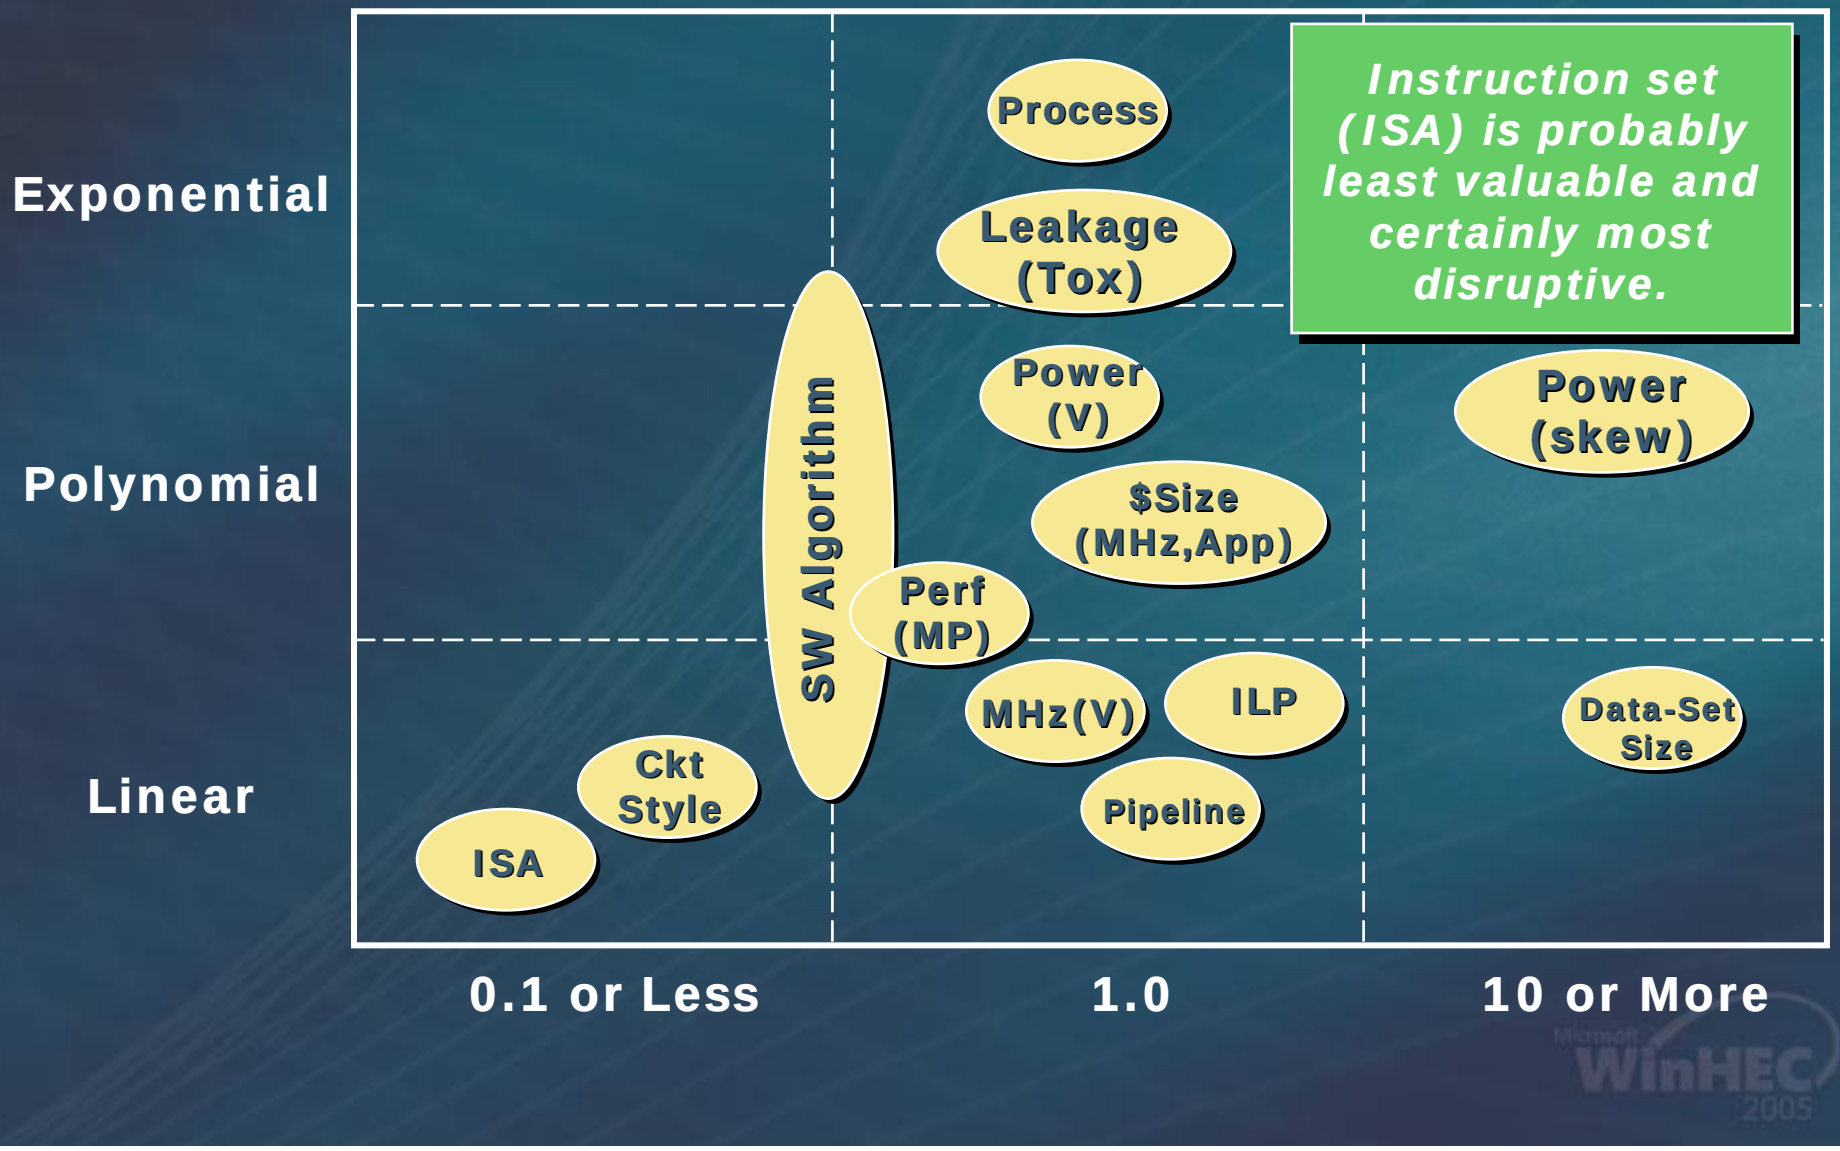

Absolute Costs Demonstrate the Time is Now

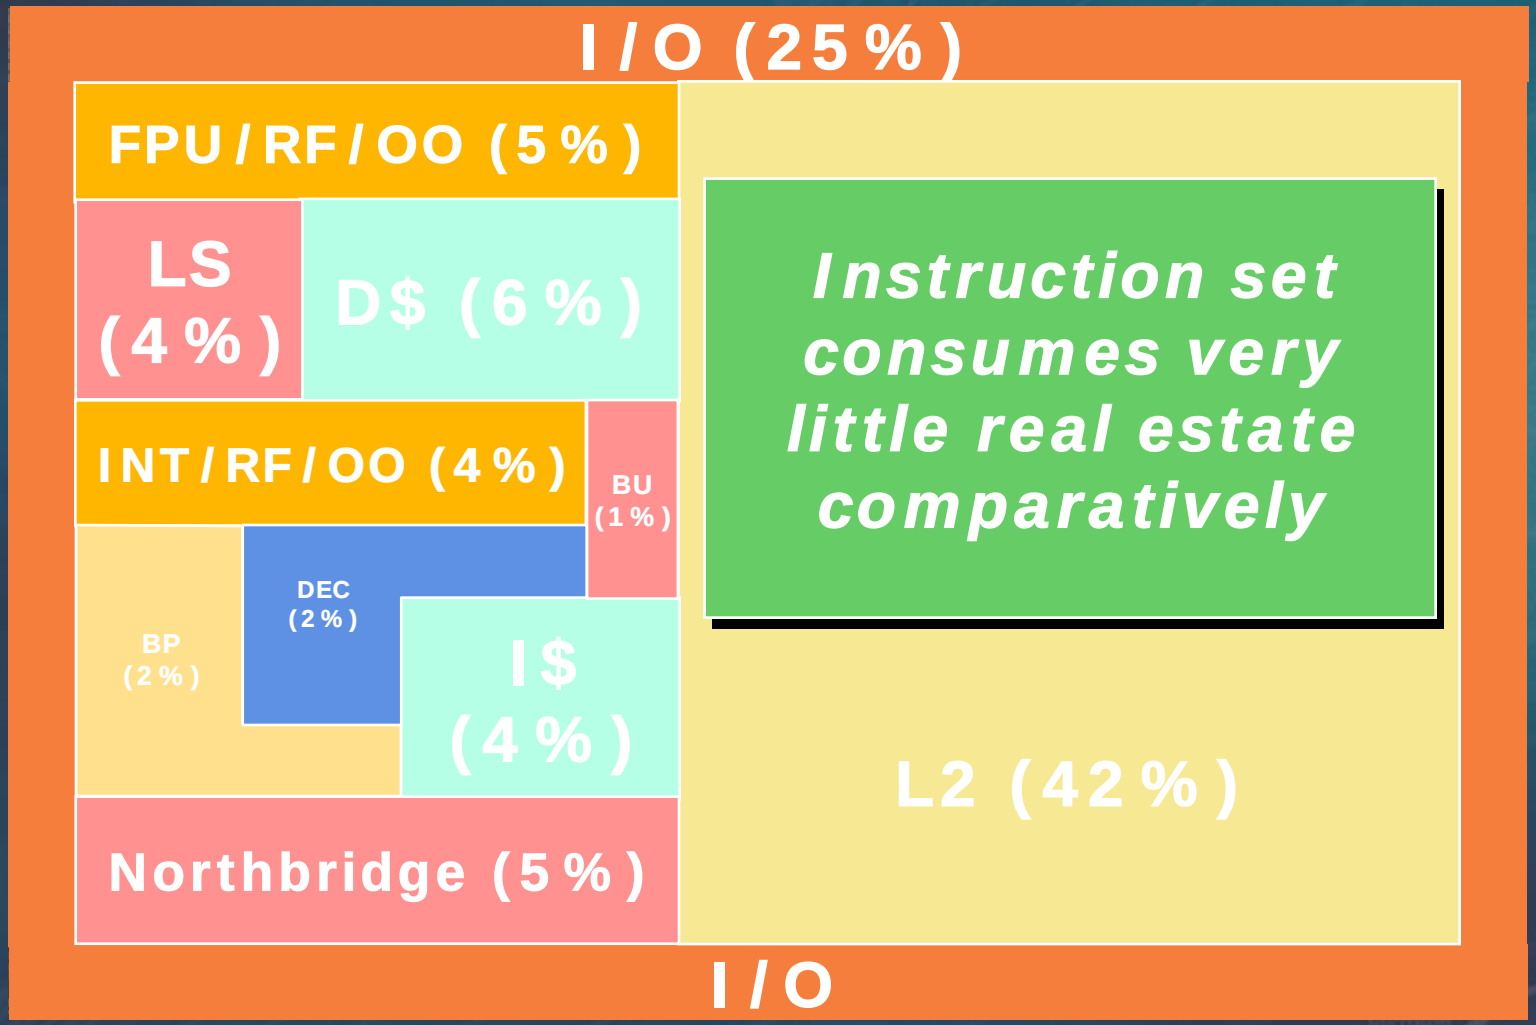

Absolute Costs Demonstrate the Time is Now

1M Transistors	1mm²
64KB Cache	1mm²
1mm²	10-20 cents
Package pin	0.5-2 cents
200 pin package	$1-$4
500 pin package	$5-$10
512KB Cache	$1-$2
Small CPU Core	5mm², $0.5-1
Substantial CPU Core	40mm², $4-$8

We Can Optimize for Area or Power or Performance.

… and so, instruction set adds very little cost.

Future Micro-Architectural Innovations

- Threaded architectures
- Multicore
- Chip level multiprocessing
- Huge scale MP machines
- Much higher performance superscalar, out of order CPU core
- Media/vector processing extensions
- Security and virtualization
- Huge caches
- Static and dynamic Power management
- Branch and memory hints
- GHz performance IO

Final Choices are driven by Optimization Priorities.

All are essentially instruction set agnostic.

Towards Instruction Set Consolidation

The trend is clear. The time is now.

Instruction Set Architecture Consolidation

Micro-Architecture Proliferation

Extending x86

Enterprise x86

Consumer x86

Low / Mid Server

SAN / NAS

Rugged Small Form Factor

Internet Appliances

LFF Consumer Electronics

High End Server

Workstation

Handheld

Desktop

SFF Consumer Electronics

Networking

Call To Action

- Hardware Developers:
  - Break through artificial barriers of power, price, form factor
  - Allow a common architecture across market boundaries
- Software Developers:
  - Do not lockout support based on ISA
  - Allow x86 to actually be Everywhere

Additional Resources

- Web Resources:
  - Specs: http://www.amd.com

    http://www.amd..com/embeddedprocessors
  - Other Resources: http://www.50x15.com
- Related Sessions
  - Low Power, small formfactor x86

questions

Microsoft®

Your potential. Our passion.™

Attribution